

MyMedia: Mobile Semantic Peer-to-Peer Video Search and Live Streaming

Matthias Klusch, Patrick Kapahnke, Xiaoqi Cao
German Research Center for Artificial Intelligence (DFKI)
Saarbruecken, Germany
{klusch,patrick.kapahnke,xiaoqi.cao}@dfki.de

Benjamin Rainer, Christian Timmerer
Alpen-Adria-University Klagenfurt, Klagenfurt, Austria
{benjamin.rainer,christian.timmerer}@itec.aau.at

Stefan Mangold
Disney Research, Zurich, Switzerland
stefan.mangold@disneyresearch.com

ABSTRACT

Mobile peer-to-peer (P2P) computing with applications such as for video on demand, file sharing, and video conferencing is gaining momentum based on new standards and technologies such as IETF PPSP, WiFi-Direct and BitTorrent live streaming. In this paper, we describe the mobile system MyMedia, that allows users to search, share and experience videos and live recordings using P2P and at the best quality possible with respect to available network capacity. In particular, the MyMedia system features a high-precision semantic P2P search and dynamic network-adaptive P2P live streaming of MPEG videos over HTTP based on the ISO/IEC standard MPEG-DASH from mobile to mobile devices in unstructured wireless P2P networks. Experimental evaluation of these features in terms of performance and energy consumption, and a first, limited evaluation of user acceptance at a film festival showed that the MyMedia system is suitable for its purpose. The MyMedia system is available as open-source software for the Android operating system.

1. INTRODUCTION

In recent years, mobile P2P computing has gained momentum based on new standards and technologies such as IETF PPSP, WiFi-Direct and BitTorrent Live Streaming, as well as an increase of investment in and support of public access to wireless broadband networks, the availability of increasingly powerful, capacitive, energy-efficient mobile devices with built-in Bluetooth and WiFi interfaces, and generous data flat-rates for users. In fact, there has been a dramatic increase in the use of mobile devices for video consumption related activities such as live streaming and shar-

ing of videos, and the user demand for relevant digital content they like will increase in the future compared to recent years [23]. In this respect, we expect that mobile P2P systems will be particularly required to feature a high-precision P2P search for video content in combination with its P2P live streaming from mobile to mobile devices. Many wireless use cases could benefit from such a combination. For example, in entertainment scenarios such as city parades, sport events, entertainment theme parks, a real time participation of visitors in local attractions and parallel events via P2P search and live streaming of relevant user-generated video content may enable an enriched user experience and improve the event logistics (e.g. by load balancing in case of underserved areas).

Though there exist quite sophisticated techniques from different technology areas for implementing these features, to the best of our knowledge, there is no mobile system with both of them. For example, semantic P2P systems [30] exploit techniques of semantic reasoning on annotated media items to significantly improve the search precision over non-semantic approaches in P2P networks with fully decentralized, i.e. unstructured, or structured DHT (Distributed Hash Table)-based overlay. Only few of these systems are for mobile environments but neither for semantic P2P search of videos nor for P2P live streaming of their search results [5, 17, 19]. On the other hand, mobile systems and applications for P2P video streaming such as Tribler Mobile [11], RapidStream [7], Peer2View [26], and the recently announced BitTorrent P2P live streaming technology [22] are addressing well known traditional concerns of integrating P2P with mobile devices such as limited battery capacities, network bandwidths and 3G traffic allowances for most users, connectivity issues related to traversing NAT (network address translator) gateways, and insufficient economic incentives for mobile app developers and content publishers like Spotify [12, 32]. However, approaches for mobile peer-assisted live streaming without centralized tracking, augmentation of distribution capacities with high-performance desktop peers acting as content boosters and computational outsourcing to public or private clouds are rarely investigated. Besides, the potential of leveraging the ISO/IEC standard MPEG-DASH

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MOBIQUITOUS 2014, December 02-05, London, Great Britain
Copyright © 2014 ICST 978-1-63190-039-6
DOI 10.4108/icst.mobiquitous.2014.258026




Figure 1: System architecture of MyMedia peer

[10, 31] for a wireless network capacity-based adaptive P2P streaming over standard HTTP remains to be unlocked. To this end, we present the first mobile system which features, in particular, a high-performance semantic P2P search and a dynamic adaptive live streaming of annotated MPEG-DASH videos from mobile to mobile devices over HTTP in wireless networks with an unstructured and semantic P2P overlay. Each of these features has been evaluated for their average performance by simulation, energy consumption on the Samsung S3 smartphone under Android, and their acceptance in a preliminary, limited user test during the 54th Thessaloniki international film festival (TIFF) which mobile application TIFF EventLive integrated these features. According to the results, the MyMedia system appears suitable and was accepted by the test users for its purpose in practice. The MyMedia system also offers peer components for dynamic semantic P2P replication and semantic service composition planning, and has been made available as open-source software for the Android operating system. The remainder of the paper is structured as follows. We provide an overview of the system architecture of a mobile MyMedia peer in section 2. The essentials of the MyMedia features for semantic P2P search and live streaming are described in sections 3 and 4 while their integration with the mobile TIFF EventLive app and its experimental evaluation are presented in sections 5 and 6, respectively. Related work is discussed in section 7 before we conclude the paper in section 8.

2. THE MYMEDIA SYSTEM

The open-source MyMedia system¹ for mobile applications under Android aims to facilitate a joint searching, sharing and experiencing of media within a mobile P2P network. In particular, MyMedia features a high-precision semantic

search of annotated media and dynamic network-adaptive live streaming of videos over standard HTTP from mobile to mobile in an unstructured P2P network. Each MyMedia peer can search, download, upload and playback media content; in its current implementation MyMedia 1.0, the media type is restricted to MPEG videos.

Mobile peer architecture. The component-based system architecture of a MyMedia peer is shown in figure 1. The semantics of MPEG videos are described in terms of semantic video services in OWL-S which are automatically generated by the peer based on the user-generated video annotation with tags from a local tag cloud, which are mapped to concepts in the peer ontology in OWL2. The peer interactively supports the user through its user interface and with its tag cloud handler in the process of semantic tagging and querying of MPEG videos. Semantic video services, metadata and video content are persistently stored and maintained on the mobile device. The peer components S2P2P [1], iSeM [14], OWLS-Xplan 1.1 [13], and DSDR [2] are implementing a high-precision semantic P2P search, selection, composition, and replication of semantic video services (cf. Sect. 3). Each MPEG video is automatically DASH-encoded/decoded by the peer for different network bandwidths with its DASH-encoder component, and displayed with its embedded VLC player component for Android. The mobile P2P live streaming of MPEG-DASH videos or live recordings within a given peer group is initiated by one of its members and then performed by all members using their pDASH components (cf. Sect. 4). The adaptive streaming from mobile to mobile devices over standard HTTP using pDASH aims at providing an optimal quality of experience with respect to available network capacities and peer resources. There is neither a central streaming server nor centralized tracking nor an augmentation of distribution capacities with additional content boosters or desktop peers, nor a computational outsourcing to public or private clouds is required for this purpose.

¹<http://sourceforge.net/projects/mymedia-peer/>

Mobile P2P interaction. The P2P interaction between MyMedia peers is concerned with IP-based wireless network communication and message routing based on the unstructured and semantic P2P overlay of the network. The network communication of peers relies on standard TCP and UDP on top of the IPv4 stack while the wireless transmission technology utilized by Android comprises WiFi (IEEE 802.11 standards), and 3G/4G communication.² Direct, bidirectional communication between MyMedia peers refers to their direct connection in the unstructured P2P overlay, and the semantic query routing bases on additional semantic overlay knowledge (cf. Sect. 3). On the network level, messages from a peer to its neighbor peers in the P2P overlay can also be routed by TCP/IP via WiFi access points depending on the underlying network infrastructure. If required, a MyMedia peer can utilize a WiFi-Direct adapter separately from its WiFi adapter in Android, on top of the used IP stack [20]. The MyMedia peer discovery which is required for an initial handshake of a peer to join the network as a new node is performed ad hoc using Bluetooth or, optionally, through a central server REST API. That is, the peer is informed about the IP addresses and TCP/UDP ports of either all direct neighbor peers by a look-up server, or just the one peer it connected to for this purpose only via Bluetooth. The initial handshake then follows a simple request/response pattern for exchanging small UDP notification messages which include IP addresses and receiving ports of the new and discovered peers. This enables the respective peers to update their local knowledge on the unstructured P2P overlay.

The MyMedia peer system has been implemented in Java for Android 3.2 with .apk size of 4 MB, and validated for different Android-based smartphones and tablets (Samsung Galaxy S2, HTC One, Samsung Galaxy Note 10.1, Sony S1, Google Nexus). In the following, we focus on the description of the search and live streaming features of MyMedia.

3. SEMANTIC P2P SEARCH

MyMedia peers perform a semantic search for relevant MPEG videos in a wireless network with unstructured and semantic P2P overlay. The used S2P2P protocol is to some extent similar to the classical random k-walker protocol for forwarding a query but driven by local knowledge on the semantic overlay rather than by random choice only.

Semantic media item descriptions. The semantic description of an item i provided by a peer p includes the unique identifiers of p, i , a text-based description of i , a (topic) term from the local tag cloud of p , a (topic) concept name with its terminological unfolding in the local ontology of p in OWL2, and/or a semantic service in OWL-S in the service registry of p for i , and the size of i . The semantic item descriptions are maintained by each peer in its local metadata repository (cf. Fig. 1). Please note that tag cloud and ontology are local to each peer, there is no global ontology in the P2P network, but all peers share a minimal common vocabulary of terms³. Though the S2P2P protocol is agnostic to the type of media items and ontology language

used for semantic annotations and computational methods for semantic relevance computing used by all peers, in MyMedia 1.0 the search is restricted to MPEG videos, semantic annotations in OWL2 and selection of semantic services in OWL-S by the selector iSeM[14].

For each video annotated by the user with tags from the local tag cloud and an additional text description, the user's MyMedia peer creates a semantic video service in OWL-S. In particular, the user-given tags are mapped to relevant concepts in the peer ontology in OWL2 based on structural ontology-based similarity measurements as proposed in [34]. These tag concepts, the textual description, and other service grounding related information are associated with respective service parameters in OWL-S for the considered video. Figure 2 shows an example of such a service in OWL-S for a video which the user described as her live recording of an actor signing the collectors edition of her latest movies. Each semantic video service is grounded in a




Figure 2: Example of semantic video service

REST service API call with standard MPEG-DASH media profile description (MPD) of the respective video (cf. Sect 4) which enables its playback by the peer with the open-source MPEG-DASH compliant VLC player for Android. In the following, we informally present the essentials of the S2P2P query routing protocol for MyMedia peers [1].

Semantic P2P query routing. A query q is issued by a peer p to k of its direct neighbour peers based on its local knowledge about their semantic expertise for q , or at random, if p has not acquired such knowledge yet. Each of these k query walkers q_k for q includes the selected topic term (tag) and concept in OWL2 (or desired semantic service description in OWL-S) for a requested item i ; TTL (time to live; number of hops of a walker) value of q_k ; requested number of copies of i ; current and suggested routing path for q_k ; the actual answer set for q_k which informs on which peer on the query path of q_k is able to provide how many (replicas) and what kind of semantically relevant items for q_k as well as their semantic descriptions are included in the piggy-backed data of q_k . Whenever a peer p' receives a query walker q_k from another peer, it adds semantic descriptions of items it knows about to the query answer set, if those items are determined by p' to be of sufficiently high semantic relevance for the query topic. The degree of semantic expertise of a peer p' for a query topic is equal to the size of its contribution to the query answer set in terms of semantically relevant items provided by p' . As mentioned above, the concrete method for semantic relevance computation is not fixed by S2P2P but is media type and application dependent. A MyMedia 1.0 peer uses its service selection component iSeM [14]

²NAT traversal is not supported by MyMedia 1.0 yet.

³These terms are the primitive components of the peer ontologies such that any peer can classify a concept it received from another peer with its logical unfolding into its own ontology without loss of semantics [30, 8]

for the hybrid semantic matching and ranking of semantic video descriptions in OWL-S with a given query. In particular, iSeM performs text similarity-based, logical concept and structural ontology-based matching, and learns the optimal aggregation of these filters in terms of precision⁴. During the search each peer observes and analyses the information of the traversing query walkers in order to learn about and update its local view of the semantic overlay, the current demands (queries) of other peers and supplies of items (semantic descriptions) by other peers in the network. The routing decision for every query walker q_k a peer p receives is based on its local knowledge of the semantic overlay. In particular, it computes a query routing path by means of a greedy path augmentation that iteratively applies Dijkstra's algorithm in order to find the shortest path with a maximal number of expert peers known to p for answering q_k while respecting its TTL restriction. Each respectively traversed peer on this path can further adjust the original query path suggested by p for q_k based on its own semantic overlay knowledge.

Eventually, each of the k query walkers returns to their original peer p by backtracking along its query path which was jointly developed by the respective peers in the network. Each of the relevant peer video links are presented to the user as soon as the respective query walker returns, or, optionally, in a top- n rank list after n query walkers returned. The subsequent selection of videos the user likes can then be followed by their peer-assisted live streaming within a given peer user group (cf. Sect. 4).

Evaluation. The experimental performance evaluation of S2P2P on a large scale has been conducted through simulation of unstructured P2P networks of one million peers with network topologies based on random power law graphs (RLPG) and random graphs (RG) with averaged connectivity 3.2. The tested item topic distributions were uniform at random and Zipf-based ($\beta = 1.05$) while ten thousand queries (TTL = 20, $k = 4$) were randomly distributed to peers. Each local peer ontology in OWL2 was generated from random samples of 20k RDF linked data items from dbpedia. The semantic relevance of data for a given query was computed by each peer based on logical concept subsumption-based and text similarity-based matching [1]. As a result, S2P2P appeared to be quite robust and offers an average precision (AP) of 0.68, which outperforms classical k-random search (AP = 0.57) as well as INGA[33] as a representative of competitive semantic P2P systems by more than 20%. If S2P2P is combined with our dynamic semantic replication strategy DSDR [2] it yields an even higher average precision of 0.82, which outperforms both k-random search with non-semantic replication [29] (AP = 0.716) and semantic replication DSDR (AP = 0.752). Domain-specific performance evaluation results for S2P2P on a much smaller scale are presented in Section 6.

4. ADAPTIVE P2P LIVE STREAMING

MyMedia peers are following the ISO/IEC standard MPEG-DASH (MPEG Dynamic Adaptive Streaming over HTTP as delivery protocol) [10, 31] for DASH-compliant video encoding and their local pDASH components to perform a network

capacity-based adaptive P2P live streaming.

DASH compliant video encoding. Each peer is using its DASHEncoder component (cf. Sect. 2) to generate MPEG-DASH compliant content by transmultiplexing and segmenting MPEG video streams from the media recorder on Android devices. In particular, a MPEG video is segmented by the publishing peer for different network bandwidths per streaming period and the respective XML-based media presentation description (MPD) file for the video is created. MPEG-DASH requires the video data to be segmented in equally sized constant time units (e.g., 1s, 2s, 4s, 6s, 10s) while the respective MPD file provides for each segment information such as different quality levels (representations), different audio streams, and subtitles. Figure 3 shows an example MPD of a video of a peer together with its current mobile device context. The *@minimumUpdatePeriod* attribute of the MPD root element denotes the update time and is used by other peers to re-download this MPD. The *@bandwidth* attribute of each video representation denotes the minimum required bandwidth for downloading one second of the video. If a peer p retrieves a video segment from

```
<MPD
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xmlns="urn:mpeg:dash:schema:mpd:2011"
  xmlns:peer="http://www.example.com/peer"
  xsi:schemaLocation="urn:mpeg:DASH:schema:MPD:2011_DASH-MPD.xsd" type="dynamic"
  minimumUpdatePeriod="PT10S" timeShiftBufferDepth="PT30M"
  availabilityStartTime="2014-07-29T10:30:02"
  minBufferTime="PT10S" profiles="urn:mpeg:dash:profile:isoff-live:2011">
  <BaseURL> http://127.0.0.1:8080/live/ </BaseURL>
  <BaseURL> http://client-peer1-IP</BaseURL> <!-- peer list -->
  <BaseURL> http://client-peer2-IP</BaseURL>

  <Period>
    <AdaptationSet mimeType="video/mp4" codecs="avc1.4D401F" frameRate="30000/1001"
      segmentAlignment="true" startWithSAP="1">
      <Representation id="v0" width="1920" height="1080" bandwidth="1000000">
        <SegmentList duration="4" initialization="liveinit.mp4">
          <SegmentURL media="live$Number$.m4s" startNumber="1"/> </SegmentList>
        </Representation> <!-- further representations for different bandwidths -->
      </AdaptationSet>
    </Period> <!-- further periods -->

  <peer:DeviceContext>
    <peer:UplinkBandwidth> 2500 </peer:UplinkBandwidth>
    <peer:PeersConnected> 2 </peer:PeersConnected>
    <peer:MaxPeers> 3 </peer:MaxPeers>
  </peer:DeviceContext>
</MPD>
```

Figure 3: Example MPD of a MPEG-DASH video

another peer p^* then the latter is adding the IP address (*@BaseURL*) of p to the peer list in the MPD of this video together with expiration time and a timestamp for each segment request of p . Besides, p^* includes its device context in the MPD which informs p about the battery level, the estimated uplink bandwidth currently available for p^* , peers connected to, and maximum number of peers allowed to connect to p^* . Each video segment request of one peer from another peer results in an update of the peer list in the MPD of the video by the latter. The DASHEncoder component can also perform an adaptive segmentation with changing segment sizes on the fly (e.g. during live recordings). For this purpose, it employs a second instance of the hardware encoder of the mobile device or utilizes FFMPEG, which allows to generate different quality representation of the currently recorded live stream. Even if the QoS parameters of the uplink do not allow peers to download the highest possible media representation, the component can trigger a transcoding to a multimedia bitrate that increases the probability of successfully delivering segments to other peers.

P2P live streaming with pDASH. As mentioned above, each MyMedia peer can download, upload and playback

⁴The selector iSeM offers a high average precision of up to 0.92 for the test collection OWLS-TC4 of the international semantic service selection contest[15].

MPEG-DASH videos on an Android device. If a peer p is requested by its user to perform a live streaming of a selected video v it has found at some peer p' to other peers in a given peer group G , then p joins G and either downloads v from the identified peer p' in order to initiate and participate in a streaming session for v in G by itself, or requests peer p' to perform the latter actions. In general, for a peer-assisted live streaming of video v from its original hosting peer p^* to other peers in G , all group members utilize their pDASH components as follows. Initially, each peer $p \in G$, $|G| = m$ which knows about p^* can retrieve the MPD of v from p^* and then start downloading the very first segment. In this case, the providing peer p^* updates its MPD of v by adding p , expiration time and timestamp of the request to the peer list of this MPD together with the device context of p^* . Second, each peer p which has downloaded video segments from p^* can provide them in turn to other peers $p' \in G - \{p, p^*\}$ on their request together with the respectively updated MPD of v at p . The updated MPDs are periodically retrieved by each peer from other peers it actually knows about in the P2P overlay. If a peer p fetches the MPD of v from some other peer p' , it extracts the actual list of peers and device contexts which are included in this MPD, and downloads the MPDs of v from these peers. Third, each peer p utilizes the device contexts in the MPDs of v it fetched from other peers in G in order to make a local decision from which of the other peers in the group to best download a maximum of n segments of v next, starting at segment s , in parallel. We implemented two alternative decision-making strategies, one basic ad-hoc and one complex optimisation-based strategy. In the first case, peer p selects a certain number of peers from the peer list stored in its MPD of the video and then tries to download segments from these target peers according to a bitrate-based adaptation logic and if its buffer fill state is high enough (e.g. 50%) [18]. In the second case, peer p determines the set $P = \{(p_i, r_{i,j}) | p_i \neq p, i \leq m-1\} \subset G$ of target peers p_i with their available representations $r_{i,j}$ of segments $j \in \{0, \dots, s+n\}$ of v by solving the following constrained optimisation (knapsack) problem [28]:

$$\max_P u(P) := \sum_{(p_i, r_{i,j}) \in P} w_j \cdot r_{i,s+j} \quad (1a)$$

$$\forall (p_i, r_{i,j}) \in P : L_u(p_i) \geq r_{i,s+j} \quad (1b)$$

$$\forall (p_i, r_{i,j}) \in P : C(p_i) + 1 \leq \text{MaxPeers} \quad (1c)$$

$$\sum_{(p_i, r_{i,j}) \in P} r_{i,s+j} \leq BW \quad (1d)$$

$$T_B - \frac{\sum_{(p_i, r_{i,j}) \in P} r_{i,s+j}}{BW} \geq \text{MinBuf} \quad (1e)$$

$$|P| \leq M_P \quad (1f)$$

In equation (1a), $u(P)$ denotes the utility function to be maximized, where w_j is the weight for a specific segment j where $\sum_j w_j = 1$. The term $r_{i,s+j}$ denotes the bandwidth (kbit/s) which is required for the $(s+j)$ -th video segment available at peer p_i (as specified in the MPD). The weighting of segments allows to prioritize certain segments for decreasing the starting time of the video content by enforcing a selection of a lower representation for the first segments. Equation (1b) requires that the upload bandwidth $L_u(p_i)$ of peer $(p_i, R_j) \in P$ (cf. Fig. 3, peer:UplinkBandwidth) shall be greater or equal to those which are specified in the requested video representation,

while equation (1c) states that the number ($C(p_i)$) of peers p_i which downloaded segments from another peer shall not exceed its maximum peer limit MaxPeers (cf. Fig. 3, peer:MaxPeers). The bandwidth needed by a representation is further constrained by the maximum available download bandwidth BW (cf. Equation (1d) and by the actual buffer fill state T_B , which shall not decrease below a threshold (MinBuf) during the download of the video segments (cf. Equation (1e)). Finally, in equation (1f) the number of segments that shall be downloaded in parallel is limited (M_P). In general, the optimization problem aims at finding a suitable set P for which the above mentioned utility function is maximized under the given constraints. After p downloads the n segments, successfully or not, the peers re-decide from which peers the next n segments shall be retrieved. This allows to incorporate new information about the peers which is periodically updated according to the *@minimumUpdatePeriod* in the MPD (cf. Figure 3). If a peer discovers changes of the device contexts of some candidate peers after having fetched their updated MPD of the video when downloading its segments, it re-computes a new feasible solution with the updated information by means of a greedy algorithm [6], which is fast but may not be optimal.

Please note that each peer applies a least-recently-used replacement strategy for deleting peers with outdated video segment requests from the peer list in its MPD of v . The number of tracked peers in a MPD is limited by their clustering into sets of peers which have approximately the same bandwidth characteristics, e.g. peers which downloaded the same representation of a video. That limits the size of a MPD and allows for using persistent connections as defined in HTTP 1.1. Besides, the number of segments for which the above optimisation problem shall be solved can be limited (e.g., 5 segments) by the number of TCP connections of a single peer and available bandwidth. Since pDASH uses the MPD as a reference to indicate the location of video segments which have already been downloaded by other peers, no structured network overlay for searching segments has to be maintained.

5. MYMEDIA APPLICATION

The MyMedia system features for recording, semantic annotation, P2P search and live streaming of MPEG-DASH videos have been integrated in the mobile application TIFF EventLive of the 54th Thessaloniki international film festival in Thessaloniki (Greece). Other features of this festival app are, for example, social ratings of films and sentiment analysis [16]. Figure 4 shows some example snapshots of the MyMedia video playlist, the tag cloud and additional text input box for video annotation, the local view of the mobile P2P network, and the video playback with the MPEG-DASH enabled VLC player for Android.

For its use at the TIFF-54 and TIFF-55 festivals in late 2013 and 2014, the MyMedia peers have been equipped with a festival specific tag cloud and ontology in OWL2 by the TIFF organizers. This default tag cloud, however, can individually evolve depending on its local usage for semantic annotation and querying of videos. General information on the TIFF festival is provided by the central TIFF server to MyMedia peers on request.

It is assumed that most TIFF EventLive users are consent with sharing video content stored on their mobiles such as live recordings from an ongoing film party, waiting queues of




Figure 4: MyMedia features of TIFF EventLive app

cinemas, or any point of interest on the festival premise in the wireless mobile P2P network at the festival. The joining of this network has been made easy for the app users to perform via Bluetooth (cf. Sect. 2). The MyMedia app has been preliminarily tested for user acceptance by a small number of selected users on the premise of the TIFF-54 festival (cf. Sect. 6).

Illustrative scenario. The following usage scenario illustrates the goals of the MyMedia system in the context of the TIFF film festival. Suppose that Sheila, Pete and Carl are befriended visitors of the TIFF each of whom being at different locations of the festival premise. They launch their MyMedia-empowered festival app on their mobiles and join the wireless P2P network for TIFF visitors. Pete uses his MyMedia peer to quickly search the mobile P2P network for latest information and media related to festival films in his preferred category, the location of cinemas where these are playing and live recordings of film parties on the festival premise. The peer returns a list of top-ranked media offered by other peers including videos of his favorite film title and a current live recording of some film party in the close proximity of the cinema where this film is playing.

Next, Pete wants to convince Sheila and Carl to jointly watch this film with him at the cinema and go to the film party afterwards. For this purpose, Pete's MyMedia peer can download both videos from the identified peers and then initiate a mobile P2P live streaming session for this content with Sheila and Carl in a peer group. Alternatively, the peer can request the identified peers to temporarily join the group for their performing of a peer-assisted streaming of their videos to Pete, Carl and Sheila. The latter is done by default in case of a P2P streaming of live recordings from the recording mobile peer.

After they have jointly watched the videos at the same time with acceptable quality and latency time, Sheila agrees with Pete's suggested plan, but Carl is hesitating. He would rather prefer a cinema which is not overly crowded and also has some restaurant in its surrounding with a nice veranda for having a social dinner before joining the film party afterwards. Fortunately, Sheila happens to be nearby the considered cinema and uses her MyMedia peer to perform and share a live recording of the waiting queue at this cinema

and a restaurant in its surrounding with Carl. This eventually leads to a joint action plan of the group for the rest of the day on the festival premise. While walking to the cinema, Pete is also consent with sharing some of his short live recordings of POIs he is passing by with other visitors and festival staff members on request for different purposes such as navigation, information, and situation-aware coordination of assistance by the festival organisers.

6. PRACTICAL EVALUATION

The practical evaluation of the MyMedia features for semantic P2P search and live streaming was concerned with (a) the performance through simulation on a small scale with 300 visitors⁵, (b) the acceptance by selected users through interviews and collected questionnaire-based feedback after their testing of the MyMedia app on the premise, and (c) the energy consumption of mobile devices on which the MyMedia-empowered festival app TIFF EventLive was tested.

The user acceptance evaluation has been conducted at the 54th Thessaloniki international film festival (TIFF-54) in September 2013 by a special evaluation group for the festival (TIFF-ELE group) which consisted of representatives of the TIFF-54 staff (6), the evaluation organizers (3), and volunteering TIFF-54 visitors (12). The TIFF-54 organizers provided one public Wifi access point on the festival premise, while the mobile devices of test users with the installed TIFF EventLive app were peer-connected through the festival WLAN (cf. Sect 2). Each MyMedia peer was equipped with a basic TIFF tag cloud of 54 concepts, and the corresponding peer ontology in OWL2. Peers providing videos were given a subset (1-3 videos) of the default set of 7 MPEG-DASH encoded videos with probability of .65, which yielded a quite sparse set of in total 390 video items in the unstructured P2P network.

P2P search performance. The search performance of MyMedia peers with S2P2P was conducted through simulation on a cluster for small-scale P2P networks of 300 peers with a power law graph-based social network topology ($\beta = 1.5$), TTL = 10, $k = 2$, video and query topic distribution following Zipf's Law ($\beta = 1.05$), and a uniform at random distribution of 7.5k queries. In particular, the MyMedia peers were physically distributed to 10 machines of the cluster, making up for 30 peers running per Java process; each of the cluster machines was equipped with 6 core Intel Xeon CPUs running at 2.67 GHz and 12 GB of RAM running Linux OS version 3.2.0-58. For these small-scale experiments, the MyMedia peers were configured to use S2P2P, in particular the service selector iSeM, but not their DSDR component for semantic replication (cf. Sect. 2). The relevance sets for test queries were set up based on test user judgments, and the average precisions (cf. Fig. 5), cumulative recall (cf. Fig. 6), and average query response times were measured.

The experimental results showed that the average precisions for peers using S2P2P were up to 1.6 times better (AP = .56) than those using traditional k-random search (AP = .35) and stabilized already after 2k queries. In fact, the MyMedia peers achieved a significantly higher search precision

⁵This limitation is based on the estimation by TIFF-54 organisation members of the average number of festival visitors who would most likely use the MyMedia features of the festival app.




Figure 5: Average precisions of S2P2P and k-random search




Figure 6: Cumulative recall of S2P2P and k-random search

with S2P2P at any time (cf. Fig.4, S2P2P AP for the latest 50 queries). Though the local search performance of each peer using its iSeM component was very high on average ($AP = .79$), the lower average precision of the whole P2P system with S2P2P is mainly due to the general problem of searching unstructured P2P networks with sparse item distributions. However, in contrast to k-random search, S2P2P was able to mitigate this problem to a large extent. Please note that the average precision of iSeM for the more general test collection OWLS-TC4 used in the S3 contest differs significantly from its local selection performance due to the domain-specific settings (different ontology, services and relevance sets) of the practical evaluation in the TIFF context. In addition, S2P2P was able to outperform the baseline in terms of cumulative recall with $CR_{10} = .32$ (versus $CR_{10} = .22$ for k-random). In the latter respect, its much more linear behavior (cf. Fig. 6) shows its ability to achieve an increase in recall even at very late steps of query routing up to the given TTL. This is mainly due to the advanced routing scheme of S2P2P where peers are making use of their local observations of semantic expertises of other peers, in particular of those which would not be reached by random walks within the TTL or due to dead ends and therefore can contribute more to a recall in late hops on average. Such knowledge is propagated by MyMedia peers on the path in the piggy-backed data of traversing queries (cf. Sect. 3). The messaging overhead of S2P2P compared to the baseline

appears very small: S2P2P just produced 1.22 times more messages. Due to their semantic expertise-driven query routing mechanism, the S2P2P peers are more likely to find and communicate with more relevant expert peers within the TTL than those performing a k-random search. The latter type of search often reached dead-ends very early, within 10 hops in the given network overlay. On average, a peer using S2P2P (baseline) processed 531 (435) queries.

The average query response time has been estimated based on the measured response times for a simple 2-peer system with Android devices (Samsung Galaxy S4, HTC One) and the observed messaging in the cluster simulation for 300 peers. As a result, the estimated average response time is 3.27 seconds with rare exceptions of up to 3.91 seconds. At the TIFF-54, we measured an average response time of about 1 second in the wireless P2P network of the festival in peer groups of two to ten test users.

P2P live streaming performance. The performance of mobile P2P live streaming with pDASH was evaluated through simulation in OMNeT++ using the INET framework with 40 peers with different bandwidths for each of them, a HTTP server with bandwidth of 1Gbit/s and 10 representations of video content with different bandwidths and resolutions. For these experiments, we have configured the peers to use a fast but naïve decision strategy rather than their alternative strategy based on constrained knapsack problem solving (cf. Sect. 4). As mentioned above, in order to select peers from which to download video segments, we configured the MyMedia 1.0 peer peer selects its target peers from the peer list stored in its MPD of the video and then tries to download segments from these peers according to a bitrate-based adaptation logic and if its buffer fill state is high enough (e.g. 50%) [18]. The experimental results showed that the bandwidth savings by peer-assisted streaming are up to 25% for the servers, and less for a representation bandwidth limit of 6 Mbit/s than for 1.4 Mbit/s since more different representations were requested by peers in this case [18]. The average latency of P2P streaming within peer groups of two to ten peers measured at the TIFF-54 festival was about 4 seconds.

User evaluation at TIFF-54. The user acceptance of the MyMedia features of the TIFF EventLive app was evaluated based on the collected responses of 18 test users (staff and visitors) from the TIFF-ELE group to a brief user questionnaire on the subject and additional personal interviews. Overall, the majority of them considered the MyMedia features very useful for enriching their personal festival experience (65% agreed, 30% neutral) and easy to use (75% agreed, 25% neutral).

In particular, the integrated search and joint live experiencing of events on the premise, POIs, waiting queues and partys shared through live recordings of other visitors or buddies was considered most important. The tagging functionality of the MyMedia peer of the festival app was accepted in general by the users (68% agreed, 33% neutral) though some of them (22%/5%) felt uncomfortable with using a tag cloud for semantic annotation/querying of relevant videos.

Most of the users (95% agreed) were positive about and satisfied with the experienced quality of the search results and displayed P2P live recordings. The measured average latency of about 1 second for P2P search and 4 seconds for P2P live streaming in peer groups of two to ten mobile de-

vices on the TIFF-54 festival premise with good wireless network conditions were considered reasonable and acceptable by all users. Few users raised concerns about the lack of integrated mobile app security technologies in the festival app for user data privacy management. A follow-up user evaluation of the MyMedia-empowered festival app is currently in preparation for the TIFF-55 festival.

Energy consumption. For measuring how much energy the P2P search and live streaming features of the MyMedia app consume on mobile devices, we used the PowerTutor2 tool, a Samsung Galaxy S3 as a test device, and the open source movie "Big Buck Bunny". In particular, we measured the energy consumed by the encoding of a live recording with the DASHEncoder, the playback of MPEG-DASH compliant multimedia content, and the semantic P2P search activity of one peer on the mobile test device.

Encoding. For all live recordings (of the movie from a 22inch LCD screen in the same environment for all experiments), we used the H.264/AVC hardware encoder, and tested resolutions of 240p, 360p, 720p and 1080p for bitrates of 500 kbit/s and 1000 kbit/s. In the following, we focus on the resolution of 240p using a bitrate of 500kbit/s (cf. Fig. 7), and the resolution of 1080p using a bitrate of 1000kbit/s (cf. Fig. 7), since these two cases showed the minimum and maximum of energy consumption in our experiments. In particular, the hardware encoder consumed less energy in the first case ($19.08J$ [Ws]) than in the second ($32.82J$), while the energy consumption of the mobile DASHEncoder itself amounts to only $3.48J$ on average in all experiments. It appears that hardware encoders on mobile devices save a lot of energy compared to software codecs [3].

Besides, the requested resolution for video display was the main factor that increased the energy consumption, accounting for $540J$ at 1080p resolution and full display brightness, while the desired bitrate had only a minor impact. These measurements resemble the results presented in [4], where display energy consumption of smart phones in various applications has been investigated.

Summing up, given the specification of the S3 battery of 2100mAh and 3.8V (7.98Wh), the process of live streaming including encoding, display of recording preview on the screen and data transmission using the WiFi adapter ($91J$) can theoretically be performed for at most 4.75 hours for the above mentioned maximum settings (4.82h for minimum settings). In practice, this result is reduced by secondary factors such as other apps running on the device and tasks performed in the background as well as physical properties related to the battery and other hardware.

Playback. Figure 9 shows the measured energy consumption during the playback of a live recorded video stream with maximum buffer size of 8 seconds. While MPEG-DASH decoding of the content consumed only $13J$, the WiFi module remained in a high energy state almost during the whole playback and consumed $91J$. Similar to the observations made for encoding, the display accounts for the largest portion of energy consumption by far. In total, the maximum expected runtime in playback mode on the S3 is around 6.29h.

Search. For evaluating the energy that is consumed by a peer on a mobile device when it searches for live videos, we set up a simple wireless mesh network with three MyMedia peers on Android devices (one Samsung S3, and two Google Nexus 7 Tablets) of which one (Google Nexus) was recording




Figure 7: Energy consumption of live recording (240p, 500kbit/s)




Figure 8: Energy consumption of live recording (1080p, 1000kbit/s)

a live stream and one (on Samsung S3) was searching for it using S2P2P (cf. Sect. 3). The amount of energy the latter peer consumed for its semantic P2P search related activities on the S3 was negligible ($0.5J$ for CPU, $0.9J$ for WiFi) compared to the one consumed for MPEG-DASH en/decoding and playback of the found live stream (cf. Fig. 10); the WiFi module was put into the *high* energy state only for a very short time period.

7. RELATED WORK

In general, the MyMedia system is most related to work on mobile P2P live streaming and mobile semantic P2P search. However, to the best of our knowledge, there is currently no system which provides both features and exploits the MPEG-DASH standard for dynamic adaptive P2P live streaming over HTTP.

For example, Tribler Mobile [11] is an open-source mobile app for Android devices for video streaming between users in a network with DHT-based overlay. The underlying Tribler system is augmenting the distribution capacities of the network with user-given content boosters and high-performance desktop computers. A video is not directly published by a peer on a mobile device to its consuming peers on mobile devices. Instead, the publisher uploads the video to all of its boosters in the network based on subscription services and then exits the app to save energy. The implemented publish-subscription mechanism relies on the use of the Twitter service for tweeting respective video links (PPSP URLs) by the publisher to Twitter subscribed consumers of a given peer group. Tribler uses decentralized DHT-based peer tracking for reverse peer discovery by which new or existing peers can periodically request DHT nodes for IP addresses and UDP




Figure 9: Energy consumption of live recording playback




Figure 10: Energy consumption of the MyMedia system during P2P search.

ports of other peers.

However, Tribler Mobile differs from MyMedia in several aspects: the Tribler system features neither a semantic P2P search for relevant videos and running live recordings, nor a P2P live streaming of live recordings from the mobile device of the recording peer to those of its consumers, nor a network capacity-based adaptive streaming based on the MPEG-DASH standard, and utilizes a structured DHT-based rather than an unstructured P2P overlay.

The Peer2View system for P2P live streaming over HTTP [26] uses a centralized tracking service for peer discovery in networks with unstructured (random) overlay, and employs a gossip-based protocol to collect QoS information such as maximum/average throughput, connectivity information, playback quality, buffering point, and network location from other peers. During P2P streaming each peer selects neighbor peers for the delivery of the stream based on this collected information, in particular their heuristically estimated upload capacities. A similar approach is taken by RapidStream [7]. However, both systems do not feature any semantic P2P video search, and MyMedia does not depend on centralized tracking for peer discovery and peer-assisted streaming of MPEG-DASH videos and live recordings.

On the other hand, as mentioned above, semantic P2P systems [30] exploit techniques of semantic reasoning on annotated media items in order to significantly improve the search precision over non-semantic approaches in P2P networks. Only a few are featuring semantic P2P search in mobile environments.

For example, the SMSN system [19] performs a semantic-based search for relevant resources and users in mobile ad hoc networks. It employs a semantic distance vector rout-

ing protocol in which peers distribute profiles of users and resources semantically described in OWL2 within a given range (k-hop limit) of their neighborhood in the network with an unstructured P2P overlay. Each peer maintains its local view on the semantic overlay in terms of which peers it knows of can provide what ontology concepts and issues semantic query walkers to them depending on the degree of semantic similarity with the given query. The open-source system has been deployed and evaluated on a set of twelve SonyEricsson K750 cellular phones and uses J2ME Bluetooth API (JSR 82) for mobile ad hoc networking.

The MP2PSW system [17] performs a search for semantically relevant items from heterogeneous sources: Web servers, P2P networks in the Web, and mobile ad hoc networks. Search in a (non-mobile) P2P network relies on Chord ring-based query routing while a mobile network is searched by exploiting IPv6 anycast, thus effectively treating the whole network as one single black box with exactly one address and unknown physical receiver peer. While this proved to be very efficient with respect to query response time, it is, in contrast to MyMedia, not actually a search in (parts of) the network, but an efficient scheme to distribute workload on coequal peers that optimally provide synchronized content. Besides, MP2PSW assumes a query agent running on non-mobile hardware in the Internet. The user may use her mobile device to issue queries and interpret results, but the actual query processing and result aggregation is performed by this central agent.

8. CONCLUSIONS

We presented the first mobile system, MyMedia 1.0, which features a high-performance semantic P2P search and a dynamic adaptive live streaming of annotated MPEG-DASH videos from mobile to mobile devices over HTTP in wireless networks with an unstructured and semantic P2P overlay. These features have been integrated in the mobile application TIFF EventLive of the 54th Thessaloniki international film festival. The experimental evaluations of the system performance and device energy consumption, and a first user evaluation at the festival revealed that the MyMedia system appears suitable and accepted by the test users for its purpose in practice. The MyMedia system is available as open-source software for the Android operating system.

Ongoing work is concerned with, for example, integrated personal recommendation by a peer based on its semantic service planning component, the follow-up user evaluation at the TIFF-55 festival, and the practical application of the MyMedia system in other domains such as film production, family travel and recreation.

Acknowledgements. This work was in part supported by the German ministry for education and research (BMBF) in project Collaborate3D under contract number 01IW1102.

9. REFERENCES

- [1] Cao, X.; Klusch, M. (2013): S2P2P: Semantic Search in Unstructured Peer-to-Peer Networks. Proc. 15th IEEE Intl. Conference on High-Performance Computing and Communications (HPCC); IEEE.
- [2] Cao, X.; Klusch, M. (2012): Dynamic Semantic Data Replication for K-Random Search in Peer-to-Peer Networks. Proc. 11th Intl. IEEE Symposium on

- Network Computing and Applications (NCA); Cambridge, USA; IEEE.
- [3] Chu-Hsing, L.; Jung-Chun L.; Chun-Wei L. (2007): Energy Analysis of Multimedia Video Decoding on Mobile Handheld Devices. Proc. Intl. Conference on Multimedia and Ubiquitous Engineering.
 - [4] Chen, X.; Chen, Y.; Ma, Z.; Fernandes F. C. A. (2013): How is energy consumed in smartphone display applications? Proc. 14th Workshop on Mobile Computing Systems and Applications (HotMobile); New York, NY, USA; ACM
 - [5] de Asis Lopez-Fuentes, F. (2010): A Routing Scheme for Content Localization in Peer-to-Peer Networks. Proc. IEEE Intl. Conf. on Electronics, Robotics and Automotive Mechanics Conference; IEEE.
 - [6] Dantzig, G. B. (1957): Discrete-Variable Extremum Problems. *Operations Research*, 5; INFORMS.
 - [7] Eittenberger, P.M.; Herbst, M.; Krieger, U.R. (2012): RapidStream: P2P streaming on Android. Proc. 19th IEEE Intl. Packet Video Workshop (PV); IEEE.
 - [8] Euzenat, J.; Shvaiko, P. (2007): Ontology Matching. Springer.
 - [9] Fitzek, F.; Charaf, H. (eds.) (2009): Mobile Peer-to-Peer. Wiley.
 - [10] ISO/IEC 23009-1 (2012): Information technology - Dynamic adaptive streaming over HTTP - Part 1: Media presentation description and segment formats.
 - [11] Jimenez, R.; et al. (2013): Tribler Mobile: P2P Video Streaming from and to Mobile Devices. In: Jimenez, R.: Distributed Peer Discovery in Large-Scale P2P Streaming Systems. PhD Thesis, KTH Royal Institute of Technology (Stockholm, Sweden), Chapter 11. Tribler Mobile (beta) at Google play, 29.1.2013.
 - [12] Kim, J.; Park, S. (2014): Resource Analysis for Mobile P2P Live Video Streaming. Ubiquitous Information Technologies and Applications, Lecture Notes in Electrical Engineering, 280, Springer
 - [13] Klusch, M.; Gerber, A. (2006): Fast Composition Planning of OWL-S Services and Application. Proc. 4th IEEE European Conference on Web Services (ECOWS); IEEE.
 - [14] Klusch, M.; Kapahnke, P. (2012): The iSeM Matchmaker: A Flexible Approach For Adaptive Hybrid Semantic Service Selection. *Web Semantics*, 15; Elsevier.
 - [15] Klusch, M. (2012): The S3 Contest: Performance Evaluation of Semantic Service Matchmakers. In: Blake, M.B.; et al. (eds.): Semantic Web Services - Advancement through Evaluation; Springer.
 - [16] Kompatsiaris, I.; et al. (2014): Real-Time Social Media Indexing and Search. Proc. Intl. Broadcasting Convention (IBC), Amsterdam, Netherlands.
 - [17] Kuo J-L; Shih, C-H.; Che, Y-C. (2013): An Efficient Query Scheme for Semantic Web on Mobile P2P Network. Proc. 5th Intl. IARIA Conference on Evolving Internet. IARIA.
 - [18] Lederer, S.; Mueller, C.; Timmerer, Ch. (2012): Towards Peer-Assisted Dynamic Adaptive Streaming over HTTP. Proc. 19th Intl. Packet Video Workshop.
 - [19] Li, J.; Wang, H.; Khan, S. (2011): A Semantics-based Approach to Large-Scale Mobile Social Networking. *Mobile Networks and Applications*, Springer
 - [20] Lombera, I. M.; Moser, L. E.; Melliar-Smith, P. M., Chuang, Y.-T. (2013): Mobile Ad-Hoc Search and Retrieval in the iTrust over Wi-Fi Direct Network. Proc. 9th Intl. Conference on Wireless and Mobile Communications.
 - [21] Loukos, F.; Karatza, H.; Kalogeraki, V. (2014): Real-time data dissemination in mobile peer-to-peer networks. *Systems and Software*, 90, Elsevier.
 - [22] Mason, M. (2014): Coming Soon: Smartphone Streaming, Powered by BitTorrent Live. BitTorrent Blog <http://blog.bittorrent.com/2014/02/14/coming-soon-smartphone-streaming-powered-by-bittorrent-live/>
 - [23] Meffert, J.; Chappuis, B.; Hazan, E.; Duncan, E.; Wager, N. (2013): iConsumers - Life online. Telecommunications, Media and Technology. McKinsey & Company.
 - [24] Mueller, C.; Lederer, S.; Pöcher, J.; Timmerer, C. (2013): libdash – An Open Source Software Library for the MPEG-DASH Standard. Proc. IEEE Intl. Conference on Multimedia and Expo (ICME); IEEE.
 - [25] Rainer, B.; Timmerer, C.; Kapahnke, P.; Klusch, M. (2014): Real-time Multimedia Streaming in Unstructured Peer-to-Peer Networks. Proc. 11th IEEE Intl. Consumer Communications and Networking Conference (CCNC), Las Vegas, USA; IEEE.
 - [26] Roverso, R.; El-Ansary, S.; Haridi, S. (2012): Peer2View: A peer-to-peer HTTP-live streaming platform. Proc. 12th IEEE Intl. Conference on Peer-to-Peer Computing (P2P); IEEE.
 - [27] Seenivasan, TV; Claypool, M (2014): CStream: Neighborhood bandwidth aggregation for better video streaming. *Multimedia Tools and Applications*, 70(1), Kluwer
 - [28] Szkaliczki, T.; Eberhard, M.; Hellwagner, H.; Szobonya, L. (2011): Knapsack Problem and Piece Picking Algorithms for Layered Video Streaming. 7th Japanese-Hungarian Symposium on Discrete Mathematics and its Applications.
 - [29] Sozio, M.; Neumann, T.; Weikum, G. (2008): Near-optimal dynamic replication in unstructured peer-to-peer networks. Proc. 27th ACM Symposium on Principles of Database Systems (PODS); ACM.
 - [30] Staab, S.; Stuckenschmidt, H. (eds.) (2006): Semantic Web and Peer-to-Peer - Decentralized Management and Exchange of Knowledge and Information. Springer.
 - [31] Stockhammer, T. (2011): Dynamic Adaptive Streaming over HTTP - Standards and Design Principles. Proc. ACM Multimedia Systems Conference.
 - [32] Sun, Y; Guo, Y.; Zhang, X.; Li, Z.; Salamatian, K. (2013): The case for P2P Mobile Video System over wireless networks: A practical study of Challenges for a mobile video provider. *Network*, 27(2), IEEE.
 - [33] Tempich, C.; Staab, S. (2006): Semantic query routing in unstructured networks using social metaphors. In: Semantic Web and Peer-to-Peer; Springer
 - [34] Wu, Z.; Palmer, M. (1994): Verb semantics and lexical selection. Proc. 32nd Annual Meeting of the Association for Computational Linguistics.